

ETL-0272

(12) LEVEL 71

Interactive digital correlation  
techniques for automatic  
compilation of elevation data

AD A109145

F. Raye Norvelle

OCTOBER 1981

DTIC FILE COPY

DTIC  
ELECTE  
DEC 3 1 1981  
S B

U.S. ARMY CORPS OF ENGINEERS  
ENGINEER TOPOGRAPHIC LABORATORIES  
FORT BELVOIR, VIRGINIA 22060

81 12 31 011


APPROVED FOR PUBLIC RELEASE. DISTRIBUTION UNLIMITED



E

T

L



Destroy this report when no longer needed.  
Do not return it to the originator.

---

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

---

The citation in this report of trade names of commercially available products does not constitute official endorsement or approval of the use of such products.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER ETL - 0272	2. GOVT ACCESSION NO. AD-A209	3. RECIPIENT'S CATALOG NUMBER 145
4. TITLE (and Subtitle) INTERACTIVE DIGITAL CORRELATION TECHNIQUES FOR AUTOMATIC COMPILATION OF ELEVATION DATA		5. TYPE OF REPORT & PERIOD COVERED Research Note
7. AUTHOR(s) F. Raye Norvelle		6. PERFORMING ORG. REPORT NUMBER
9. PERFORMING ORGANIZATION NAME AND ADDRESS U.S. Army Engineer Topographic Laboratories Fort Belvoir, VA 22060		8. CONTRACT OR GRANT NUMBER(s)
11. CONTROLLING OFFICE NAME AND ADDRESS U.S. Army Engineer Topographic Laboratories Fort Belvoir, VA 22060		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 4A762707A855, B, 0026
13. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		12. REPORT DATE October 1981
		13. NUMBER OF PAGES 41
		15. SECURITY CLASS. (of this report) Unclassified
		18a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for Public Release; Distribution Unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Anaglyphic Stereo Display on Video Equipment Digital Correlation Interactive Stereo Image Matching		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) A program using digital correlation techniques has been implemented on the Digital Image Analysis Laboratory (DIAL) to extract elevation data automatically from stereoscopic pairs of digital images. The DIAL enables the operator to view on a display unit the stereoscopic model and the computed match-point coordinates in 3-D using anaglyphic techniques. The match points are presented as 3-D profiles superimposed over the 3-D terrain model, and consequently, the operator can observe the compilation process, detect errors as soon as they occur, and interact with the process to prevent the accumulation of error or to tune the parameters as conditions warrant. The program is described in terms of the hardware, the algorithm, the operating procedures, and the examples of compiled areas.		

DD FORM 1473

EDITION OF 1 NOV 55 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

403292

elt

The work reported on was done under DA Project 4A762707A855,  
**PREFACE** Task B, Work Unit 0026.

The work was performed under the technical guidance of Mr. Michael A. Crombie and the supervision of Mr. D.E. Howell, Chief, Information Sciences Division; and Mr. L.A. Gambino, Director, Computer Sciences Laboratory.

COL Daniel L. Lycan, CE and COL Edward K. Wintz, CE were Commanders and Directors and Mr. Robert P. Macchia was Technical Director of the Engineer Topographic Laboratories during the study period.

Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
<b>A</b>	

## CONTENTS

TITLE	PAGE
PREFACE	1
ILLUSTRATIONS	3
INTRODUCTION	4
THE DIAL	6
THE IMAGES	9
THE ALGORITHM	10
General	10
Window and Search Area Sizes	11
Point Spacing	12
Window Shaping	13
Match Point Prediction	16
Correlation Coefficients	17
Quality Tests	18
Iterative Techniques	19
Match Point Adjustment	21
OPERATING PROCEDURES	21
Image Display	21
Floating Dot	22
Parameter Selection	22
Starting Profiles	23
Automatic Compilation	23
Data Inspection and Editing	25
EXPERIMENTAL RESULTS	26
General	26
Match Point Adjustment Experiment	27
Window Shaping Experiment	30
Number of Iterations Experiment	31
Match Point Prediction Experiment	34
Data Editing Experiment	35
Correlation Coefficient Experiment	37
DISCUSSION	38
General	38
Viewing Problems	39
Execution Time	40
CONCLUSIONS	41

## ILLUSTRATIONS

FIGURE	TITLE	PAGE
1	The Interactive Subsystem of the ETL Computer Facility	6
2	The DIAL Work Station	7
3	RECO1 Image of the Phoenix Area	9
4	Relationship Between the Window and the Search Area	12
5	Match Point Spacing on the Left and Right Images	13
6	Relationship Between the Shaped Window and the Search Area	14
7	Geometric Basis for Window Shaping and Prediction of "Next" Point	15
8	Iterative Software Logic	20
9	Match Point Adjustment Scheme	21
10	Compilation of Bad Area Without Match Point Adjustment	27
11	Compilation of Bad Area With Match Point Adjustment	29
12	Compilation Standard Without Match Point Adjustment	29
13	Compilation Without Shaping and Without Match Point Adjustment	30
14	Compilation With One Iteration and Without Match Point Adjustment	32
15	Compilation With One Iteration and With Match Point Adjustment	33
16	Compilation With No Match Point Prediction and Without Match Point Adjustment	34
17	Stereogram of Match Point Lines Before and After Editing	35

## INTERACTIVE DIGITAL CORRELATION TECHNIQUES FOR AUTOMATIC COMPILATION OF ELEVATION DATA

### INTRODUCTION

It is within the purview of the Computer Sciences Laboratory (CSL) of the Engineer Topographic Laboratories (ETL) to apply digital image manipulation techniques to the solution of various mapping, charting, and photointerpretation problems. Over the past 10 years, CSL has sponsored in-house and private industry efforts to develop digital techniques that can be applied to extract elevation data automatically from a stereoscopic pair of photographs. One of the first efforts was performed for ETL by Keuffel and Esser Company (K&E)<sup>1,2</sup> and was aimed at developing and testing an algorithm for digital correlation using the "infiltration method," a complex process that follows the path of easiest correlation rather than a defined path in local space. Much of the philosophy developed by the K&E effort was used later in a comprehensive, five-phase study by Control Data Corporation (CDC)<sup>3</sup> aimed at developing algorithms, implementing the algorithms on special purpose, high-speed "Flexible Processors," actual testing on stereoscopic imagery, and developing preliminary specifications for a digital stereo image- processing system.

Prior to and overlapping the CDC effort, in-depth studies were performed in-house by Crombie<sup>4,5,6</sup> to isolate the cause-effect relationships between digital correlation results and changes to the various combinations of parameters that can be used in a digital stereoscopic image matching process. For example, the studies involved tests of several statistical measures of correlation to determine which give the better

---

<sup>1</sup>P. Rosenberg and K.E. Erickson, *Digital Mapping System Study; Final Report*, U.S. Army Engineer Topographic Laboratories, Fort Belvoir, VA., 1971.

<sup>2</sup>P. Rosenberg, K.E. Erickson, and G.C. Rowe, *Digital Mapping System: Mathematical Processing*, ETL-CR-74-6, U.S. Army Engineer Topographic Laboratories, Fort Belvoir, VA., May 1974.

<sup>3</sup>D.J. Panton, M.E. Murphy, and D.S. Hanson, *Digital Cartographic Study and Benchmark; Final Report*, ETL-0168, U.S. Army Engineer Topographic Laboratories, Fort Belvoir, VA., 1978.

<sup>4</sup>M. Crombie, P. Lem, and T. Hay, *Single Photo Analysis of Sampled Aerial Imagery*, ETL-RN-74-10, U.S. Army Engineer Topographic Laboratories, Fort Belvoir, VA., August 1974, AD-A012 176.

<sup>5</sup>M. Crombie, *Semiautomatic Pass Point Determination Using Digital Techniques*, ETL-0051, U.S. Army Engineer Topographic Laboratories, Fort Belvoir, VA., December 1975, AD-A026 082.

<sup>6</sup>M. Crombie, *Stereo Analysis of a Specific Digital Model Sampled from Aerial Imagery*, ETL-0072, U.S. Army Engineer Topographic Laboratories, Fort Belvoir, VA., September 1976, AD-A033 567.

results under various conditions. Tests were also made to determine what size the patches should be in the correlation process, what the spacing should be between correlation points, what is the effect of patch shaping, etc. In the process of conducting the tests, using digital correlation techniques, software was generated that eventually evolved to a complete package capable of the automatic compilation of corresponding photo coordinates from a stereoscopic pair of photographs. This program was run, as were the others, in a batch mode, and consequently the results had to be evaluated using laborious, time-consuming statistical analyses of large amounts of data.

In order to examine, evaluate, and edit the data from the in-house studies more readily, ETL decided to generate the appropriate software for the Digital Image Analysis Laboratory (DIAL), which ETL was developing at the same time that the CDC and in-house studies were being conducted. With the advent of the DIAL, which will be described later in this report, a means was available for image display and manipulation along with the opportunity to interact with the software controlling the computational process. This effort, aimed at the visual examination of batch mode results, soon evolved into a complete DIAL compilation program, which hereafter will be referred to as the Digital Interactive Mapping Program (DIMP).

The DIMP was generated to serve the following purposes:

1. To minimize the time required to perform digital correlation tests and to evaluate results.
2. To perform studies for algorithm development.
3. To test means for visual display of results and for data editing.
4. To demonstrate an automatic, but interactive, means to compile elevation data, using digital image processing techniques.

In implementing DIMP, several features suggested by the CDC studies were incorporated, especially regarding the stereoscopic presentation of the digital images, the stereoscopic presentation of the correlation results, and the techniques of data editing. On the other hand, the results obtained in the in-house efforts were relied on heavily in implementing an algorithm and in selecting options and parameters, such as window sizes, point spacing, quality checks, etc.



The DIAL makes it possible to select and display the stereoscopic model in three dimensions (3-D) using anaglyphic techniques and, at the same time, to view the correlation results as 3-D profiles superimposed over the stereoscopic model. Errors in the automatic, continuous compilation process can be detected visually (as opposed to statistically) as soon as they occur, and the operator can interrupt and interact with the software to get the process back on track. The operator can also select specific terrain features to compile and can determine immediately what effects changes to various parameters have on the correlation process. Just how these features have been implemented will be discussed in the following sections by describing the DIAL, the DIMP algorithm, the operating procedures, and some experimental results.

**THE DIAL** The Digital Image Analysis Laboratory (DIAL) is a large-scale, interactive system designed for research and development activities in digital image processing. It is part of the ETL computer facility, which consists primarily of a host computer subsystem, a STARAN subsystem, and an interactive subsystem. The equipment that comprises the interactive subsystem is given in figure 1.

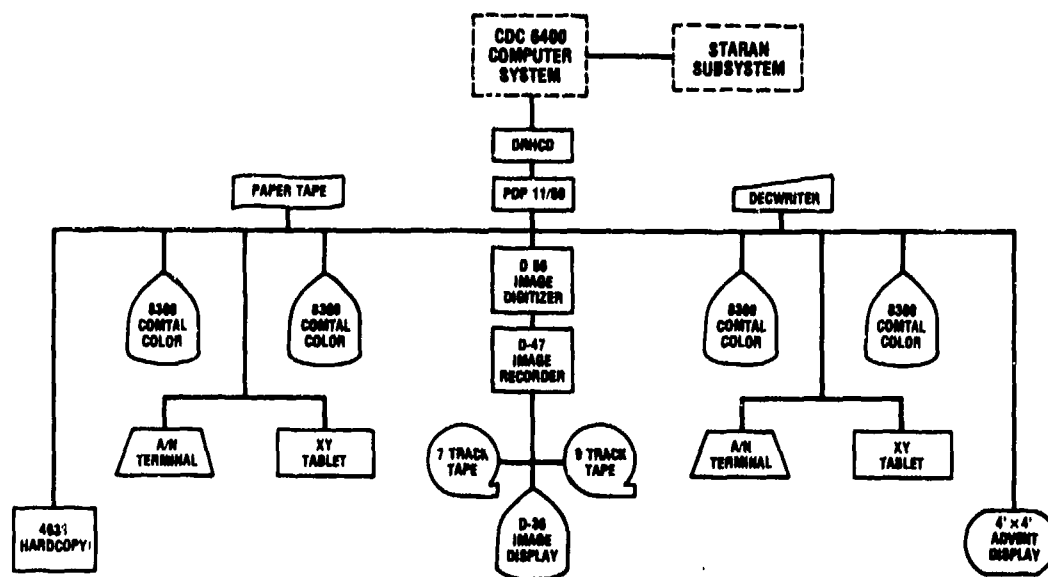


FIGURE 1. The Interactive Subsystem of the ETL Computer Facility.

The Control Data Corporation CDC-6400 computer is used in a time-sharing mode with other ETL users as the host computer for the DIAL system, and thereby performs practically all of the processing and control functions. The Goodyear Aerospace STARAN computer is an associative array processor that is used, because of its high-speed and parallel-processing capabilities, to perform computationally bound tasks. Because it is not used with the DIMP program, it will not be discussed further. The interactive subsystem of the DIAL is that part which makes the DIMP possible since it enables the operator to display and view the digital photographic images and the correlation results and, when deemed necessary, to interrupt the processor to change program parameters and variables.

The interactive subsystem consists mainly of two work stations, one of which is shown in figure 2. Both work stations are controlled by a DEC PDP 11/50 computer that is interfaced to the CDC-6400. The host computer does practically all of the processing, but some simple local operations can be done concurrently on the PDP 11/50. Each of the two work stations consists of two COMTAL 8300 display monitors, a TEKTRONIX alphanumeric terminal, an XY-tablet, and one trackball unit per monitor. One TEKTRONIX hardcopy device is used to copy information from both terminals.



**FIGURE 2. The DIAL Work Station.**

Each display unit features refresh storage of three digital image memories characterized by an array of 512 x 512 elements (or pixels) whose gray shades can vary in magnitude from 0 to 255 (8-bits). Digital images larger than 512 x 512 pixels, or whose gray shades are larger than 8 bits, are resampled by the system software to reduce the size for display purposes, but the full image remains available for computational purposes via disk storage. Each monitor also has three 1-bit, 512- x 512-pixel overlay memories and one 1-bit, 512- x 512-pixel fast-graphics memory. The graphic overlays are used to superimpose graphic images, such as lines, over the 8-bit images in color. Since three overlays are available, three different color overlays may be present at any one time.

The fast-graphics plane is used to plot graphic symbols, such as dots, lines, circles, etc. over the image, giving a fourth color graphic capability. The monitors can be operated in a black and white or true-color mode at the discretion of the operator. In the black and white mode, one of the three image memories is displayed on the monitor. In the true-color mode, any combination of the image memories is displayed in red, blue, and green to give a color picture. In the DIMP, the true-color mode is used to display the left image of a stereoscopic pair in both blue and green to give cyan. The right image is displayed in red. By using the anaglyphic glasses, the operator can view the superimposed red and cyan images in 3-D. A 3-D profile can be viewed by plotting one line in cyan and the other in red on the graphic overlays, and likewise a "floating" dot can be obtained by using a red cursor and an equivalent size dot plotted in cyan on the fast graphics plane. The profile and dot applications are described later in more detail.

The cursor on each display can be defined as various symbols and in various colors. The trackball unit is used to change the position of the cursor on the screen, and "select" and "done" buttons on the trackball unit are used as a signal to the computer to record the cursor coordinates for subsequent use in the executing program. Likewise, the program can command the PDP 11/50 to position the cursor on the monitor to a computed position. The XY-tablet functions similarly to the trackball. It is not discussed here because it is not used in the DIMP.

The TEKTRONIX alphanumeric terminal is used by the executing program to send questions to and receive answers from the operator. For example, at the beginning of the program, the executing program asks the operator to enter the file name of the image to be displayed. Execution does not resume until the operator types in the file name on the keyboard of the alphanumeric terminal. Another use of the terminal involves the "interrupt" button. The DIMP is programmed to check periodically whether or not the operator has pressed the "interrupt" button on the keyboard. If it has been pressed, the program loops to a section of the program that prints out a menu on the terminal screen. At this point the operator has the opportunity to select that part of the executing program that enables the operator to choose a new subimage, change the operating parameters, print out statistics gathered thus far, etc.

Practically all experiments that were run with the DIMP  
**THE IMAGES** involved a stereoscopic pair of images called RECO1, shown in  
figure 3, and RECO2. The images are small subsets taken from  
photographs of the Phoenix, Arizona area. The photographs were alined and digitized  
on a microdensitometer comparator with the direction of flight parallel to the scanning  
axis of the comparator. Since the original photographs are at the same scale and  
essentially vertical, there is no Y-parallax, and practically no skew exists between  
images. The images were digitized at a pixel and line spacing of 24  $\mu\text{m}$  (micrometers)  
using a spot size of 34.5  $\mu\text{m}$ . A total of 2048 pixels per line and 2048 lines were  
recorded with a 10-bit representation of each gray shade value. The digitized data  
were recorded on tape and, as required by the DIAL, transferred to a disk pack.



**FIGURE 3. RECO1 Image of the Phoenix Area.**

The original photographs were taken from an altitude of approximately 23,500 feet with a 6-inch focal length camera. The scales are therefore about 1:47,000, and each pixel covers about 3.8 feet on the ground. Since the DIAL displays can display only 512- x 512-pixel images, a full RECO1 or RECO2 is resampled to give a four-times reduction by displaying only every fourth pixel and line. The resolution is thereby reduced by four, and consequently, subsets are selected from the 2048- x 2048-pixel images in order to operate with full resolution. Also, the original 10-bit gray shades were remapped to 8 bits to conform to the COMTAL constraints.

The digitized imagery represents about 2 x 2 inches of area on the original 9- x 9-inch photographs. When displayed on the 10-inch screen of the monitor, the viewing scale is approximately five times larger than the original, or 1:9400. When a subset of 512 x 512 is taken from the 2048 x 2048 image, the subset is displayed at a viewing scale of 1:2350.

The terrain covered by the digital images consists of various types of features, such as an orchard, flat ground with numerous shallow drains, extremely steep terrain with good content, and relatively steep and rolling terrain with very little content suitable for correlation. There is little detail in the form of towns, buildings, or dense networks of roads.

The original photographs have been triangulated previously, giving all the data necessary to describe the orientation of the photographs relative to an absolute ground coordinate system. This is not significant, relative to the DIMP, but these data were used off-line in another program to compute X, Y, and Z-ground coordinates based on match point data derived by the DIMP.

**General** • Given the interior and exterior orientation data and the calibration data for a stereoscopic pair of photographs, only the corresponding X and Y photo-coordinates of an image on the stereoscopic pair are needed to compute its ground position. The digital correlation algorithm described herein is used to obtain the necessary corresponding photocordinates.

**THE ALGORITHM**

In this scheme, a window of gray shade values centered around the point in question on the left photograph is compared at every possible position within a specified search area to the gray shades on the right photograph. The position within the search area where the maximum correlation occurs will define, to a pixel precision, the location on the right photograph of the corresponding image. A curve is fitted in both the X and Y directions to the maximum correlation value with ones on each side, and it is then analyzed to determine to a fraction of a pixel where the peak occurs. The specified pixel coordinates of the image on the left photograph and the computed pixel position on the right photograph can be transformed off-line to actual photocoordinates and subsequently used to compute X-, Y-, and Z-ground coordinates.

Some of the individual considerations that must be made to implement the continuous and automatic compilation of match (or corresponding) points are discussed in the following sections.

**Window and Search Area Sizes** • As a matter of definition, the term "window" refers to that rectangular array of gray shades on the left photograph that is symmetrical relative to the pixel position for which a match point is required. The "search" area is a rectangular array of gray shades on the right photograph symmetrically centered about that pixel position where the match point is estimated to occur. The estimated point in the search area is refined by the correlation process to give a computed match point.

The window size used in the DIMP is generally  $15 \times 15$  pixels, although the operator has the option of selecting any practical size and shape of window. The search area is generally two lines and six pixels ( $17 \times 21$ ) larger than the window, but any size can be used. These sizes are large enough to give good results, yet small enough to allow reasonable computational speed. A  $15 \times 15$  window passed over a  $17 \times 21$  search area results in three rows of correlation values in the Y-parallax direction and seven columns of values in the X-parallax direction (figure 4). This allows for plus or minus one pixel of Y-parallax removal and plus or minus three pixels correction to the estimated X position of the match point.




FIGURE 4. Relationship Between the Window and the Search Area.

The window and search area sizes are initially selected by the operator, but they will be altered under program control when the program logic detects conditions that call for larger or smaller array sizes. For example, if previous match point quality indicates that the process is "lost," a larger window and/or search area will be used to acquire better correlation results over a larger local area. If previous results indicate that the process is not lost, but the quality of the new point is suspect, the window and search area will be reduced in order to find good correlation in a smaller area closer to the predicted match point. This control helps to prevent the selection of a false peak in an area where multiple correlation peaks occur.

**Point Spacing** • The compilation process proceeds from left to right and from top to bottom at equal, operator-selected increments (represented by NCRE) on the left photograph. The spacing on the right photograph will be varied because of image displacement, caused primarily by relief displacements. Since match points are determined in a regular array in image space, a post-processing scheme must be used to transform any subsequent elevation data into an even array in ground space.

Figure 5 illustrates a case where a match point is determined on the top starting line of imagery at every, say, fifth pixel on the left photograph. The process is then repeated for every fifth line of imagery. Any point spacing, NCRE, can be used in the DIMP, but as described in the following section, options such as window shaping depend on the point spacing. If the point spacing is too large in comparison with the window size, window shaping cannot be performed accurately. A five- or seven-pixel spacing has been used most frequently in the DIMP because such spacings are consistent with other available options and provide acceptable speeds of operation.




FIGURE 5. Match Point Spacing on the Left and Right Images.

**Window Shaping** • Since stereophotos are taken from different perspectives, a length of sloping terrain imaged on the left photograph will be different in size and shape from its image on the right photograph. To make the two images more nearly alike, one must reshape one or the other. Since match point computations are performed in image space, only one image needs to be shaped relative to the other. If ground space computations were used, both images would have to be shaped and unnecessary computation time would result.






FIGURE 6. Relationship Between the Shaped Window and the Search Area.

As illustrated in figure 6, the smaller window of the left photo is shaped to match the search area on the right photograph by using the inverse equations in (1):

$$Y = Y_T$$

$$X = (X_T - B \cdot Y_T) / A \quad (1)$$

where

$X_T, Y_T$  = New pixel coordinates

$X, Y$  = Original pixel coordinates

$A, B$  = Transformation parameters

The shaping equations compute the positions in the original unshaped array of gray shades that correspond to the desired whole number pixel coordinates of the shaped array. The  $X$  and  $Y$  values normally will be fractional values of pixels, and consequently, a bilinear interpolation of surrounding gray shades is used to compute gray shade values for the shaped array.

The transformation parameters A and B are computed on the basis of previous match point results. In the general case, the previous four columns to the left of the "next" point and the previous four rows are used in the computation. In figure 7, the DX values represent the distances between match points on the right photograph in the direction of X-parallax. A weighted average of the DX values is used to compute the scale, A, relative to the incremental point spacing, NCRE, on the left photograph:

$$DX_{12} = (DX_1*1 + DX_2*2 + \dots, DX_{10}*10 + DX_{11}*11)/66$$

$$A = DX_{12}/NCRE \quad (2)$$

Likewise, the skew between images, B, is computed as follows:

$$DY_{12} = (DY_1*1 + DY_2*2 + \dots, DY_{10}*10 + DY_{11}*11)/66$$

$$B = DY_{12}/NCRE \quad (3)$$




FIGURE 7. Geometric Basis for Window Shaping and Prediction of "Next" Point.

The weighted computations are used to give the previous, closer match point results more weight, but yet include further points to minimize the effect of errors in the immediate vicinity of the "next" point. The A and B computations are based on only four columns and four rows of previous results in order to minimize computer storage requirements. Modified versions of equations (2) and (3) are used when less than four rows and four columns have been compiled, but the approach is generally the same.

**Match Point Prediction** • The success of the DIMP depends largely on how accurately the match point location on the right photograph can be predicted. The correlation process is, after all, a local refinement to an estimated location and not a global search operation. The Y coordinate of the first point in each row (I, 1) is determined by incrementing the Y coordinate of the match point above (I - 1, 1) by the point spacing, NCRE. This assumes that the two digital images are at the same scale. The Y coordinate of each point after the first is considered to be the same as that of the previous point (I, J - 1).

The X coordinate of the "next" point is determined using the weighted distance, DX12, as described in the previous section. As shown in figure 7, the DX12 value is added to the X value of the previous point in the same row (I, J - 1) to give one estimate for the next point. A second estimate for the next point is obtained by projecting a straight line through the three previous points (I - 3, J; I - 2, J; I - 1, J) in the same column of match points. The two estimates are averaged, and the result is used for the X coordinate of the next point (I, J).

At the beginning of the correlation process, modifications to the above techniques are used to compute the location of the next point until four rows and four columns of match points have been determined. For the first row of points, and the first point of each row, the process uses the X and Y coordinates determined by the operator in the "starting profiles" procedure, which is described in a later section. For rows 2 and 3, the X values are computed using only the first estimate for the point based on the DX values available thus far.

**Correlation Coefficients** • The operator has the option of using the linear correlation coefficient (RXY), the covariance correlation coefficient (SXY), or the absolute difference coefficient (DXY) to determine the degree of correlation between the window and the search area. The three coefficients are computed as follows:

$$RXY = \frac{\sigma_{xy}}{\sigma_x \sigma_y} = \frac{\Sigma (X - \bar{X})(Y - \bar{Y})}{\sqrt{\Sigma (X - \bar{X})^2} \cdot \sqrt{\Sigma (Y - \bar{Y})^2}}$$

$$SXY = \sigma_{xy} = \frac{1}{N} \Sigma (X - \bar{X})(Y - \bar{Y})$$

$$DXY = \frac{1}{N} \Sigma \left| (X - \bar{X}) - (Y - \bar{Y}) \right|$$

where

X, Y = the gray shade values for the window and search area.

$\bar{X}, \bar{Y}$  = the average X and Y.

N = number of elements in the window.

The differences between the gray shades and the mean ( $X - \bar{X}, Y - \bar{Y}$ ) are used in the computations rather than the absolute values (X, Y) to negate the influence of brightness differences between the images. The SXY and DXY values are still influenced by contrast differences, but RXY is not. The RXY coefficient is normalized and will always take on an absolute value ranging between 0 and 1, but SXY and DXY values will vary in range. It is important to note that in cases where both X and Y images have constant gray levels (such as a lake, snow, etc), the RXY and SXY will indicate poor correlation results, since their values will be very small if not indeterminate. On the other hand, DXY will give a zero value, which indicates a high degree of correlation. In the case of adverse areas, then, DXY can be misleading in the absence of other correlation measures, such as the sharpness of the correlation curve.

The primary advantage of using SXY or DXY is that they require fewer computations and therefore are faster and cheaper to use than RXY. This is due to the need with RXY to compute standard deviations,  $\sigma_x$  and  $\sigma_y$ , which are not

used with SXY and DXY. For a given correlation computation,  $\sigma_x$  is computed only once for the window, whereas  $\sigma_y$  is computed for the search area numerous times. For example, a 15- x 15-pixel window can be correlated in 21 different positions within a 17- x 21-pixel search area, and therefore, a  $\sigma_y$  would be computed for twenty-one 15- x 15-pixel subsets taken from the search area. In an attempt to normalize somewhat the value of SXY, the covariance  $\sigma_{xy}$  was divided by  $\sigma_x^2$  (in some tests) in order to keep SXY approximately between 0 and 1. Note that in the ideal case,  $\sigma_x^2 = \sigma_x \sigma_y$ ; therefore, the modified SXY computation becomes identical to RXY. Although this technique helped control the size of SXY, the range still varied widely in the presence of contrast differences. Because the RXY value is more consistent and predictable, it was used almost exclusively in the DIMP, even though it takes longer and costs more to compute it.

**Quality Tests** • The quality of correlation is judged good or bad, based on the magnitude of the computed correlation coefficient and also on the magnitude of correction the correlation process makes to the predicted location of the match point. The terms "good" and "bad" are used here to indicate conformance to statistical tests and not as an absolute measure of quality. The larger values of RXY and SXY and the smaller value of DXY indicate the higher degree of correlation. In the DIMP experiments, absolute values of RXY below 0.35 were considered to be bad. This threshold was determined more or less arbitrarily, based on experiments with the images. No maximum value was used with DXY, nor minimum with SXY, since they are not normalized and their ranges could not be predicted accurately. As found by Crombie,<sup>7</sup> the RXY is the most dependable of the three and was used almost exclusively in the DIMP experiments, even though it takes longer to compute.

Assuming a suitable correlation coefficient is obtained between the window and search area, one must apply a test to check how much correction will be applied to the predicted match point. If the correction exceeds a certain limit set by the operator, the point is considered to be bad, and the predicted position rather than the computed position is retained as the match point coordinates. Normally, a correction limit of two pixels is used with the DIMP, but smaller or larger values will be substituted under program control, depending on the history of the previous match point quality.

---

<sup>7</sup> M.Crombie, *Stereo Analysis of a Specific Digital Model Sampled from Aerial Imagery*, ETL-0072, U.S. Army Engineer Topographic Laboratories, Fort Belvoir, VA, September 1976, AD-A033 567.

When the correlation process succeeds in passing the quality test, a value of zero is assigned to the quality of that point. If it is not successful, a value of 1 is added to the quality value of the previous point in the same column of match points. If, for example, the correlation process has failed five times along a vertical column of match points, the quality of the fifth point is 5. The operator can set a limit on how many successive times the correlation process can fail. When the limit is reached, the automatic process halts and signals the operator to intervene, if necessary, in order to get the process back on track. The operator can also intervene at will in order to circumvent the accumulation of bad points in obviously poor correlation areas. Just how the operator gets the correlation process back on track is described later in the section on operating procedures.

**Iterative Techniques** • To enhance the chances for acceptable correlation results, the operator can allow a multiple number of iterations on each point. The logic used in the iterative technique is given in figure 8. For example, if a maximum of four iterations is allowed, The DIMP will first perform the computations based on the initially selected window and search area and the maximum allowable correction to the predicted location of the match point. If the correlation value is too small, the window and search areas are increased in hopes that a larger area will result in a larger correlation value. If a suitable correlation value is obtained before four iterations, the logic then checks to see if the correction to be applied to the predicted match point is within the allowable tolerance. If it is, the point is declared good, and the process continues at the next point. If the correction is too large, then the logic checks the quality of the previous point ( $I - 1, J$ ). If the quality was bad, the logic assumes that the process is lost, and a larger correction is needed.

More iterations will be performed with the same size window, but with larger search areas and larger permissible corrections, until either the quality tests are passed or, because the process reached four iterations, the point is considered to be bad. If the quality of the previous point is good but the correction is too large, the logic assumes that the process is not lost, but false correlation peaks occur for some reason. More iterations are then performed with increased window and search area sizes, with the allowable correction remaining constant. In this case, a larger local correlation area should determine a stronger peak correlation value within acceptable limits of the predicted point. In a special case when an acceptable correlation value is obtained on the first try (the previous point was good, but the correction is too large), it is assumed that multiple correlation peaks occurred with a false maximum. The window and search area sizes are then reduced for the next iteration in hopes of finding a suitable correlation value in a more local area relative to the predicted match point.

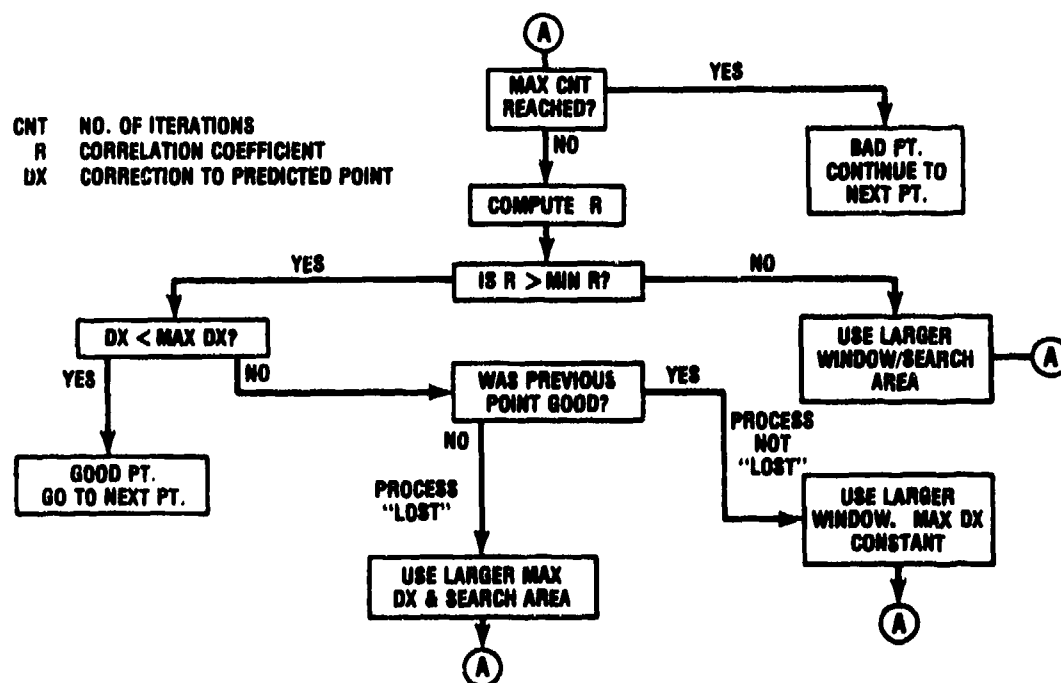


FIGURE 8. Iterative Software Logic.

**Match Point Adjustment** • After the DIMP completes the correlation process of all points in a given row, the quality of each point in the row is checked to determine whether any were judged to be bad. If bad points are encountered, a straight line interpolation is done between the good points on each side of the bad ones to compute an adjusted match point coordinate for the bad points (figure 9). The interpolation scheme takes advantage of the fact that, for the most part, the terrain is continuous and uniform over short distances. The interpolated values are considered to be better than the "bad" points, but for the purpose of the iterative logic described in the previous section, the points are still judged to be bad.




FIGURE 9. Match Point Adjustment Scheme.

#### OPERATING PROCEDURES

**Image Display** • The initial step in operating the DIMP is to enter through the keyboard of the TEKTRONIX alphanumeric terminal the names of the two photographs that constitute the stereoscopic pair. The computer reads the digital data from a disk pack and makes the images available for display on one of the display units. As is the case of RECO1 and RECO2, images larger than 512-x 512-pixels are reduced in scale so that the full image will fit on the screen. The left photograph is displayed first, and the operator uses the cursor to define the top-left corner of a 512- x 512-pixel subset of the image that is to be compiled. The right photograph is then displayed, and the operator selects the same general subset as selected from the left photograph. The two subsets are read from the disk and then displayed on the second monitor at full resolution. With the use of the color capability of the monitor, the left subimage is displayed in cyan and the right subset in red. The operator then uses anaglyphic glasses to view the subset imagery as a 3-D stereoscopic model.



In selecting the subsets, the operator can remove the Y-parallax and gross X-parallax displacements caused by possible differences in scanning and digitizing the photographs. The two images must, however, be scanned parallel to each other and must be at nominally equal scales. Otherwise, a preprocessing scheme must be used to put the imagery in a format suitable for stereoscopic viewing. Likewise, non-vertical imagery must be transformed, and properly scaled, before it is used in this manner.

**Floating Dot** • The color and fast-graphic capabilities of the monitors make it possible to superimpose a "floating dot" in the 3-D model. The floating dot is constructed in the DIMP by defining the cursor as a 3- x 3-pixel array whose color is red. In conjunction with the cursor, a 3- x 3-pixel array is plotted on the fast-graphics plane of the monitor. The trackball associated with one monitor is used to move the cursor dot and the cyan dot equally throughout the stereoscopic model. The second trackball is used to move the cyan dot in the X- and Y-parallax directions relative to the red cursor. With the anaglyphic glasses on, the operator sees a white floating dot, which can be moved throughout the model area and, by adjusting the parallax between the red and cyan dots, can be "raised" or "lowered" in elevation to appear to rest on the 3-D surface. The floating dot appears to rest precisely on the terrain surface when the red and cyan dots are over corresponding images on the red and cyan images, respectively. When this occurs, the operator, by definition, has determined manually the match point on the right photograph for the image position defined by the dot on the left photograph. Consequently, the operator can do manually what the DIMP does automatically with digital correlation techniques.

**Parameter Selection** • Once the full resolution subset of imagery to be compiled has been selected and displayed, the operator is prompted by the computer to enter the values for the parameters discussed in the previous sections on "The Algorithm." When executing the program, the operator can stop the program and change these parameters to suit the terrain conditions better, but normally this is not necessary.

**Starting Profiles** • The operator has the option of compiling as much or as little of the 512- x 512-pixel area as desired. Large areas would normally be selected, but in special cases, such as recompilation of bad areas or the study of select terrain features, small areas are desired. In any case, the operator selects the beginning point at the top left corner of the desired area by placing the floating dot precisely on the 3-D surface at that point and recording its coordinates on the left and right images by pushing a button on the trackball unit. The operator then proceeds from left to right to plot a profile of the terrain surface. In this process, the Y coordinate on the left photo remains constant so that the profile will track precisely along a specific line of imagery on the left photo. At every major break in the terrain surface, the operator sets the dot on the ground and pushes the trackball button to record the match point coordinates.

After the desired areas have been traversed from left to right, the procedure is repeated from top to bottom, starting at the previously defined beginning point on the left edge of the compilation area. The DIMP then interpolates between the gross profile points to compute "given" estimated match point coordinates at every increment along the profile that corresponds to the previously selected point spacing between correlation points. The interpolated match point coordinates are then used as the "predicted match points" described previously and are also used for window shaping of the first row in the correlation process.

**Automatic Compilation** • After the operator completes the process of drawing the starting profiles in the stereoscopic model, the correlation process becomes automatic until the operator intervenes or the program stops due to loss of correlation. Starting at the first row, the correlation process proceeds left to right across the area designated by the operator when the starting profiles were drawn. At the conclusion of each match point determination, the floating dot is positioned at that match point until the next match point is determined. The operator, therefore, observes the floating dot moving from left to right and "up" and "down" as it follows the shape of the terrain surface. Actually, the dot is not usually positioned at each new match point, but rather at some increment selected by the operator. For example, if correlation is acquired at every fifth pixel, the operator may specify that the dot be repositioned for

every fourth correlation - or at every 20 pixels. The reason for this is that for the closer spacings, the dot would move too fast for the operator to view in 3-D. By observing the moving dot, the operator can tell if the correlation is going properly, which also indicates to the operator that the program is working and that no equipment failures have occurred.

At the completion of each row of correlation points, lines are drawn columnwise between the last two completed rows of match points. On the left image the lines are vertical and evenly spaced, since the points are predesignated when the operator selects the incremental spacing between correlation points. These lines are drawn in cyan, using the fast graphics plane of the COMTAL display, and are superimposed over the cyan terrain image. Red lines are drawn columnwise on a graphic overlay between the last two completed rows of computed match points and are superimposed over the red terrain image. These lines are neither straight nor evenly spaced, because the match point coordinates reflect the X-parallax between images, mainly because of changes in elevation. Since the cyan and red lines pass through corresponding images, the operator perceives 3-D profile lines that follow the terrain surface. If the correlation results are in error, the lines will appear to "float" above or below the terrain model surface. The operator uses this visual display of the correlation results to determine if and when manual intervention is needed to prevent the propagation of errors in the compilation process.

The lines are plotted at the completion of each row of correlations, and as a result, the lengths of the lines grow progressively with time. This type of display enables the operator to study casually the previous results obtained in the stereoscopic model and also to anticipate just what problems may be encountered because of upcoming changes in the terrain surface. If the correlation process proceeds from steep to flat terrain, for example, the operator may elect to stop the program and speed up operation by inhibiting those options not needed in flat areas. If no operator intervention is necessary, the automatic mode continues until the area selected by the operator is completed. Before proceeding to compile a new subimage, the operator will normally study the visual results on the display and will decide whether any areas of the stereoscopic model were not compiled satisfactorily. In the presence of bad results, the operator can correct the data before continuing with new areas or can elect to correct the data at a later date, possibly off-line from the DIMP.

In the current version of the DIMP, on-line data editing consists of recompiling the bad areas. This process differs from the normal process only in that the starting profiles drawn by the operator are very close to the bad correlation area, and therefore the influences of errors in previous match points are negated. The recompilation process works extremely well and minimizes the need for postediting-type schemes so often associated with automatic compilation systems. On the other hand, the problems that caused bad results sometimes prevent suitable recompilation, and alternative editing techniques must be used. Some of these techniques are discussed in the following section.

**Data Inspection and Editing** • The DIMP algorithm and the associated interactive features are designed to minimize the occurrence and the further propagation of errors. The program automatically allows several iterations on each individual point in order to achieve suitable correlation. If this fails, interpolated – and usually better – match point coordinates are computed at the completion of each row for each point judged by the software to be bad. When lines are drawn through the match points, the operator then has the opportunity to intervene on the next row to prevent further propagation of errors. If the operator does not catch the errors, then the program will stop after reaching a preset maximum number of consecutive bad correlations. However, by the time the operator intervenes or the program stops, it is too late to correct errors. Rather, the operator can only keep matters from getting worse.

There are times when the above steps are not totally sufficient to prohibit the accumulation of “patches” of errors that must be corrected. The global view of the stereoscopic model, along with the plotted profiles, enables the operator to inspect the results before proceeding to the next area to be compiled. This is the most convenient time to recompile the bad patches because the errors are easily detected, no new setups are required, the recompiled areas are very small, and the process goes very fast. The recompilation process can succeed where the original failed because new starting profiles are drawn close to the bad areas, thereby negating the influence of the previous bad match points used in point prediction and window shaping. Some image areas, such as those containing scratches, watermarks, and other artifacts, defy recompilation efforts because the correlation results tend to come out the same regardless of the approach. In these cases, an alternate data-editing scheme must be used.

In one such scheme, the operator manually constructs the 3-D profile lines in those areas where automatic correlation fails. The operator does this by the same procedure that was used to construct the starting profiles. That is, he sets the floating dot on the ground at detectable breaks in the slope of the terrain and records the match point coordinates given by the pixel coordinates of the cursor dot. The software then interpolates values for points in between to fill in the column of match point coordinates at the desired incremental spacing. These values are then substituted for the erroneous ones. This technique is not part of DIMP, but is used instead in an off-line data editing program. It is relatively time-consuming when there are several profile lines to correct, but it works quite well in many cases. In some cases, however, the operator, because of poor scene content, cannot correlate poor image areas any better than the automatic process, and therefore both the recompilation and the manual correlation techniques fail.

A third data-editing technique can be used with good results under all the conditions described. It is one in which the bad correlation area is outlined by the operator; then the software adjusts the bad match point values to conform to the shape of the terrain as defined by the surrounding good match points. A simple linear interpolation between good points would usually suffice, since these error patches are normally small in area. More sophisticated curve-fitting and terrain-modeling techniques could be employed in special, complicated areas. The techniques have not been used in the DIMP per se, but have been applied at ETL to smooth the elevation data. The success of such techniques depends on the assumption that the terrain surface is continuous and uniform over small areas.

#### **EXPERIMENTAL RESULTS**

**General** • One of the more desirable aspects of the DIMP is that the correlation results are displayed visually. Consequently, the operator can judge the quality of the results quickly, without needing laborious statistical analyses. The operator views the results with anaglyphic glasses, but for documentation purposes the cyan and red lines and images can be displayed in black and white, photographed, and made into stereograms suitable for publication purposes. The stereograms are used herein to explain some of the results obtained with the DIMP.

The results were obtained with a specific stereopair of images, which are not necessarily global in application. The results are given mostly to help emphasize and clarify the previous narrative on the options allowed by the algorithm and the concept of the visual display of results.

**Match Point Adjustment Experiment** • It has been pointed out that after completing each row of correlations, points that were judged "bad" are assigned new match point coordinates based on interpolation between "good" points on each side of the bad ones. Very often the "bad" points are actually good, but simply did not pass the minimum quality tests. In a few cases, "good" points are actually bad. Nevertheless, the interpolation process rarely causes worse errors and, in many cases, is downright necessary. One such case is shown in figure 10.



**FIGURE 10. Compilation of Bad Area Without Match Point Adjustment.**

Figure 10 shows the results of compilation in an area where lack of terrain detail, rapidly changing directions of slope, and steep terrain made correlation extremely difficult. In this case, match point adjustment was not applied as an option, and the match points are plotted wherever they occur. One can see that the results are entirely useless, since in most areas the profiles bear no resemblance to the actual terrain surface and most profiles are so far off that they cannot be viewed stereoscopically. Correlations were determined at every seven pixels and every seven lines using a 15- x 15-pixel window and a 17- x 21-pixel search area. Every third possible profile line was drawn, and the lines are therefore spaced 21 pixels apart. A maximum allowable correction to the predicted match point was two pixels, and up to a maximum of five consecutive bad points were allowed before the program would stop and ask for operator assistance. The compiled area contains 1477 correlation points, of which 367 were judged to be bad. Five consecutive bad points were encountered 39 times and required that the operator set the dot back on the ground to get the process back on track. Otherwise, the results would have been even worse. Of the 1477 points correlated, 999 required corrections between zero and two pixels, and 111 required over two pixels.

The results shown in figure 11 were obtained using interpolated values between good points in lieu of the correlated bad points. Out of 1477 points, only 167 (versus 367 previously) failed to pass the quality tests, and the program required operator intervention three times rather than 39 as in the prior case. A total of 1205 (versus 999) points required corrections between zero and two pixels, and 105 required more than two pixels. The reason that the interpolation technique works better than no interpolation is that it provides better match point coordinates for "bad" points than does the point prediction scheme. As a consequence, the window shaping and point prediction for the "next" point are reasonably accurate, and the chances for the correlation process to correct itself are enhanced greatly. In this particular experiment, the interpolation process makes the difference between achieving relatively good results and receiving absolutely worthless results. The area shown in figure 10 and figure 11 is the poorest on which to correlate in the whole of RECO1 and RECO2. In the better areas, interpolation is not as significant, because the correlation process does not get lost so easily and does not need the added advantage of interpolation to find or maintain good correlation after encountering bad previous match point results.



**FIGURE 11. Compilation of Bad Area With Match Point Adjustment.**

An example of a good area is shown in figure 12. This area was compiled using the same parameters as before and without adjustment to the bad points through interpolation. This scene is shown here to illustrate a case in which interpolation is unnecessary; but it will be used again in subsequent discussion as the standard against which other results are compared. It is necessary that the standard for comparison does not include adjustment to match points through interpolation, because this process is so successful that it sometimes makes poor parameter and option selections look good and therefore leads to erroneous conclusions. This will be illustrated later in the section on "Number of Iterations Experiment."



**FIGURE 12. Compilation Standard Without Match Point Adjustment.**



**Window Shaping Experiment** • The "standard for comparison" shown in figure 12 was compiled with a 15 x 15 window, a 17 x 21 search area, a two-pixel maximum correction to the predicted point, eight consecutive losses before stopping, and four possible iterations per point. Window shaping was performed, the point prediction option was applied, the linear correlation coefficient was used, and no interpolation was performed. In figure 13, the same area is compiled with the same parameters and options as the standard except that no window shaping was performed. A close inspection with a pocket stereoscope shows that, because of poorer correlation accuracy, the profiles in figure 13 generally are not as smooth as those in figure 12. In the standard case, 159 out of 3265 points required corrections to the predicted match points of more than three pixels, and 28 points were judged to be bad. With no shaping, 210 points required corrections of more than three pixels, and 91 points were bad. Corrections greater than two pixels indicate that the process has to search, rather than refine, because of poor correlation results. Also, the average correlation coefficient is higher with shaping than without (.80 versus .76), indicating better similarity between images when they are shaped. The correlation results shown in figure 13 would have been much worse. However, the operator intervened five times and "set the dot on the ground" to minimize the accumulation of additional error. No intervention was required with the standard case.



**FIGURE 13. Compilation Without Shaping and Without Match Point Adjustment.**

Shaping is not done in any case if the ground is flat; accordingly, the differences between shaping and non-shaping experiments with the DIMP will be obvious only in the rougher terrain. Other shaping versus non-shaping experiments were conducted using various areas in RECO1 and RECO2 which contained rough terrain. Shaping invariably improved correlation but sometimes not enough to justify the added computational time. In other cases, as in this one, it made the difference between getting excellent results and losing the process. Except for flat areas, the operator will not know whether shaping is necessary, and therefore would include this option in all runs.

**Number of Iterations Experiment** • In the case of the "standard" results shown in figure 12, 3180 out of 3265 points, or 97.4 percent, were correlated successfully on the first try. Thirty-two points needed two tries, 16 needed three tries, and 37 needed four. Only 9 out of 37 passed on the fourth try, leaving only 28 bad points - a success rate of 99.1 percent. It seems reasonable to accept the 97.4 percent obtained with one iteration and call the others bad rather than spend the additional and disproportionate computational time on multiple iterations trying to recover the extra 2.6 percent of the points. To demonstrate the significance of the additional 2.6 percent of the points, an experiment was run using the same parameters and options as the standard except that only one try per correlation point was allowed. That is, the original window and search areas were correlated for each point, and if the correlation value was too low or the indicated correction to the match point was too large, the point was judged to be bad and the predicted match point coordinates were used in lieu of the computed values. The process would then continue with the next point instead of iterating on the bad point in hopes of obtaining good correlation.

The results obtained with only one try and with no match point adjustment through interpolation are shown in figure 14. Compared to the standard results, this experiment gave poor results, and they would have been even worse had not the operator interrupted the correlation process 15 times to set the dot on the ground and get the process back on track. Whereas 97.4 percent of the points passed on the first try when four tries were allowed, now only 92.3 percent passed in a single try. In other words, the attempt in the standard case to recover 2.6 percent of the points through multiple iterations accounts for a total increase in successful correlations of 6.8 percent (99.1 minus 92.3) of the points.



**FIGURE 14. Compilation With One Iteration and Without Match Point Adjustment.**

A success rate of 92.3 percent in a single try is, on the surface, quite good. However, because the bad points were grouped together, the correlation process was lost. The operator had to intervene to prevent the accumulation of worthless data, and in so doing, kept the apparent success rate quite high. Still, the results were unacceptable because the operator had to interact with the program too often and would have had to correct too many of the profiles (figure 14) by post-editing or recompilation.

The reason why one try at successful correlation often fails is that when a bad point is encountered, the match point coordinates assigned to that point are those that were predicted. When the next point in that column is tried in the next row, correlation may be bad again because the new predicted point, being based on the previous bad results, is possibly in error. With multiple iterations, the previous history of bad points would allow the correlation process to perform some searching for the match point as opposed to a small refinement to a predicted point. Usually, the search process is successful in reestablishing good correlation, and bad points are prevented from propagating. Since only one try was allowed in this experiment, searching was not permitted, and once the correlation process was lost, it remained lost until the operator intervened. This string of bad points shows up in the stereogram as profiles "floating" above or below the terrain surface.

As pointed out earlier in the match point adjustment experiment, the poor selection of parameters and options used in the DIMP are sometimes made to look good when adjustments are made to bad points by interpolation between good points. This happens because the new adjusted coordinates are better than the predicted match point coordinates that are assigned as the coordinates when the match quality is bad. As a consequence, point prediction for the next point and window shaping are more accurate, and the chances for success on the next point are enhanced significantly. The influence of the bad points acquired because of poor parameter or option selection does not bear so heavily therefore on the success of the future results. This point was illustrated in figure 11 and again in figure 15.



**FIGURE 15. Compilation With One Iteration and With Match Point Adjustment.**

The results shown in figure 15 were obtained in the same manner as those in figure 14 except that now interpolation is used between good points to replace the match point coordinates of the bad points. Whereas 250 points were bad when no interpolation was used, now only 65 are bad, and the results are practically as good as those in the standard case (figure 12).

**Match Point Prediction Experiment** • Inasmuch as the correlation process used in the DIMP provides a refinement to a predicted position of the match point, it is imperative that a predicted point be computed correctly in order that the amount of refinement or correction stays within the maximum limit specified by the operator. To demonstrate the value of point prediction as used in the standard case, an experiment was run in which no point prediction was used. Instead, the X coordinate of the previous point (I - 1, J) was used for the predicted point. The results of the experiment are shown in figure 16. Although there are numerous errors in the profiles, there would have been more except that the operator intervened eight times to set the dot on the ground, thereby preventing the accumulation of additional strings of errors. The total number of bad points was 138 as compared to 28 for the standard case, and 233 (versus 85) points required multiple iterations in order to achieve suitable correlation. Also, more points required larger corrections in this case than in the standard case.



**FIGURE 16. Compilation With No Match Point Prediction and Without Match Point Adjustment.**

Various areas in the stereoscopic model were compiled with and without point prediction, and generally results were significantly better with point prediction. This was especially true in the areas where the profiles changed direction, such as when they crossed a ridge line running parallel to the X-parallax direction. Without reliance on the point prediction scheme to predict the slope change based on the previous results,

the profile line would simply be projected straight ahead, and when viewed stereoscopically, would appear to cantilever out from the top of the ridge line. If the terrain content is poor, the correlation process will fail, and the bad results will simply propagate until the operator intervenes. In the flat areas, the slope of the profiles does not change appreciably, and the lack of point prediction has very little significance in the correlation process.

**Data Editing Experiment.** • Figure 17 is a stereogram that shows the difference in the match point profiles before and after data editing. Both the left and right terrain scenes are taken from the same photograph (RECO2) and therefore appear to be flat. The left scene contains a plot of the match points before data editing, and the right scene, after data editing. Although the terrain appears to be flat, differences in the match point lines before and after editing will appear in 3-D. The reader can therefore view the stereogram with a pocket stereoscope and determine where data editing was performed and estimate its value.



**FIGURE 17. Stereogram of Match Point Lines  
Before and After Editing.**

Several patches of error were corrected in the left portion of figure 17, the most significant being around the black dot located to the upper left of the trail. This black area appears on RECO2, but not on RECO1, and probably is caused by some type of damage to the photograph. At any rate, the correlation process failed in this area, and consequently, the match point lines appear jagged rather than smooth like the terrain surface. The fact that the lines are jagged indicates that the correlation process found "good" points here when they were actually bad. Otherwise, the match point adjustment process would have assigned new values to the bad points, and they would have been plotted smoothly across the artifact. The mismatched points in the blank area affected the results further along in the process, giving rise to bad results slightly into the good area appearing below the black dot. In this instance, the option to allow multiple iterations per correlation point works to the detriment of the DIMP. If allowed to search in the presence of artifacts, the correlation process will sometimes lock onto false correlation peaks that just happen to meet the minimum correlation coefficient requirement. Since the process is in a search mode, the allowable correction to the predicted match point is increased and the chances of satisfying all statistical requirements are enhanced. As a consequence, correlation on artifacts could give marginal, but statistically acceptable, results.

Since it was obvious to the operator that the results around the black dot indicated, by the absence of match point adjustment, that false correlation had occurred, the area was recompiled using tighter parameter control. For example, the minimum acceptable correlation coefficient was raised from .35 to .50. The maximum allowable correction to the predicted match point was set to one pixel rather than two pixels, and only one iteration, rather than four, was allowed for each point. With tighter constraints, more points were judged to be bad, and the interpolation process for the bad points produced more realistic and smoother match lines across the black dot. In areas where there were errors but no artifacts, the error patches were simply recompiled using the same parameters as the original compilation; but the starting profiles were drawn near the error patch, thereby giving reliable match point prediction and window shaping data that were not available for some reason in the original case.

**Correlation Coefficient Experiment** • Several runs were made using various areas of RECO1 and RECO2 to demonstrate the difference in results when RXY, SXY, and DXY are used to measure the quality of correlation. No stereograms of the results are shown, because the results varied so unpredictably that to show results of one or two cases would be misleading. In all cases, the RXY coefficient gave good results. In areas easily correlated because of excellent scene content, all three coefficients worked well. In those areas where scene content was very poor, the RXY coefficient was more consistently good and SXY and DXY fluctuated between good and terrible.

The correlation coefficient serves two purposes in the DIMP. First, it is used to detect the position within a small array of where the maximum correlation occurs. Second, its value is tested to determine the quality of correlation. To test the correlation value, one must know before hand what minimum value is acceptable for RXY and SXY and what maximum value is acceptable for DXY. In the case of RXY, the absolute value of the coefficient always ranges between 0 and 1, but SXY and DXY are not similarly constrained. As a consequence, the acceptable limits on SXY and DXY cannot be specified reliably before the compilation. In cases in which the test limit on the correlation coefficient was, by luck, set properly by the operator, the SXY and DXY coefficients worked well within limits. Too often, however, the test limit was too loose, and all points were judged good when many were actually bad. When the test limit was too tight, not enough good points were obtained to keep the process on track. In either case, the correlation process simply failed.

Some tests were made that essentially eliminated the magnitude of the correlation coefficient as a quality measure. For example, the minimum acceptable RXY coefficient was set to zero, and therefore all points could pass this test. It remained, however, that the maximum correction to the predicted point had to stay within specified limits. In these tests, the process worked well until points were encountered that required multiple iterations. Too often, poor but - in this test - acceptable correlation results were obtained, and by coincidence the maximum correlation value occurred within acceptable limits from the predicted point. Consequently, the point was judged to be good when actually it was bad, and the errors propagated from then on. This happens mostly with SXY and DXY but, to a lesser extent, with RXY also. The consistency and predictability of RXY are the reasons RXY was used in all the experiments described herein, in spite of the fact that it is more time consuming to compute.



**General •** There are several ways in which a suitable algorithm can be implemented, and the DIMP is just one of several. The DIMP serves the purpose for which it was developed, and it is an excellent tool for conducting tests for the purpose of algorithm development, simulation of other compilation schemes, and trade-off analyses. The main attribute of the DIMP is that it is interactive; therefore, the operator can anticipate when and where interaction is necessary and prevent the propagation of errors. This is made possible because the DIMP provides a global view of the stereoscopic model and of the correlation data in the form of 3-D profile lines. Another important feature is that the operator can choose small, discrete areas for compilation, and test the effects of specific terrain features on the correlation process. This was important in the early stages of the DIMP development when appropriate logic had to be designed to handle the unique correlation problems imposed by features such as orchards. Test data derived from orchard areas eventually led to a logic scheme that worked well in both the orchard and general areas.

As far as the DIMP algorithm is concerned, two features can be singled out as having major significance. First, the iterative capability was found to be a powerful means for maintaining good correlation results in adverse correlation areas. This technique essentially changes the DIMP philosophy, when required, from one of refinement to a predicted match point to one of a limited search operation over a small area. It was found, generally, that 90 percent of the time a refinement was sufficient but that the search operation kept the process on track in the poor areas, thereby allowing the refinement process to remain appropriate in the good areas. In other words, without the search logic, the process would get lost in the bad areas, remain lost in the good areas, and thereby negate the use of a refinement technique. Second, one has the option of replacing the match point coordinates of bad points with those derived by linear interpolation between good points. This is done at the completion of a row of correlations, and therefore the advantages are available immediately to correct bad data in that row and enhance the chances for suitable correlations in the next row via more accurate match point prediction and window shaping.

Although the DIMP serves the purposes for which it was intended, there are some weaknesses in the operation, dealing mainly with viewing problems and speed of operation. Some of these weaknesses are described in the following sections.

**Viewing Problems** • When the graphic planes of the display monitors are used to plot symbols, such as lines, the terrain imagery displayed on the screen is replaced at the involved pixel locations by the graphic overlay data. Consequently, if a red line is viewed through the anaglyphic glasses, the cyan lens filters out the red line, and the operator sees a black line instead. Likewise, the red lens filters out the cyan line, and a black line is seen by the operator. When the match point profiles are viewed stereoscopically, then the operator sees one white, 3-D profile line, and one monoscopic black line on each side of the white one. When the separation between the red and cyan lines is large enough, the black lines are separated so that they are not objectionable and do not interfere with stereoscopic perception. Very often, the lines will be separated adequately at the lower elevations, but they begin to cross each other at the higher elevations. This does not present much of a problem, because usually the interference occurs only over short lengths of the profile. In the flat areas, both the cyan and red lines are essentially straight and parallel to each other. However, it sometimes happens that most of the cyan lines precisely overlap the red lines, and stereo perception is lost. The solution here is to select the subimages with more or less X-parallax bias so that the red and cyan lines are staggered and do not interfere with each other.

A second viewing problem occurs with the graphic planes when the red and cyan lines (or dots) have to occupy the same pixel position. Since this is not possible, only one color is seen, and depth perception of the dot or line is impossible. This type of problem occurs when, for example, the non-straight red lines cross the straight cyan lines. The 3-D perception of the dot is lost in cases where the stereo imagery appears to be in the plane of the display screen. That is to say, there is not enough X-parallax bias between the two subimages that comprise the stereopair, and consequently the red and cyan dots will overlap. Although the viewing problems are objectionable at first, they are relatively infrequent, and the operator soon gets used to the situation. One solution to all of these problems is not to use the graphic overlays to display the dots and lines but rather to modify the gray shades of the red and cyan terrain images to incorporate the line and dot symbols. This, of course, is not a viable solution, because an inordinate amount of computer time would be required to modify gray shades continually and rewrite the updated imagery to the COMTAL display. Execution time is already a problem, as will be discussed in the following section.

**Execution Time** • In the case of correlation with a 15 x 15 window and a 17 x 21 search area, the DIMP operates at a rate of about 9 correlations per second. The rate decreases when multiple iterations are involved and increases when smaller window or search areas are used. This rate does not appear to be low to an operator, but compared to hardwired, automatic compilation equipment, the DIMP is very slow. Approximately 80 percent of the time spent on each point is used by the subroutine that computes the linear correlation coefficient. Obviously, then the way to speed up the program is to use specially designed hardware to perform the correlation coefficient computations. As stated in the ETL studies performed by Control Data Corporation,<sup>8</sup> special hardware can provide a 34 to 1 improvement in execution time compared to the general purpose CDC 6400 computer.

It was stated earlier that the STARAN associative array processor was a subsystem of the ETL computer complex, but it was not discussed because it is not used with the DIMP. Actually, one test was made using the STARAN computer to perform the correlation computation in the DIMP, and it provided an eight-times increase in speed compared to the CDC 6400. Although the STARAN is capable of much faster computation, the overhead cost of manipulating data in the CDC 6400 before it was sent to, and after it was received from, the STARAN was the limiting factor and greatly reduced the effectiveness of the STARAN. The correlation routine on the STARAN currently is fixed to handle only a 17 x 17 window correlated over a 23 x 23 search area or a 23 x 23 window over a 31 x 31 search area. Neglecting input/output and waiting on the CDC 6400, the STARAN could perform about 140 correlations per second with the smaller arrays. The DIMP, however, relies heavily on an iterative approach with varying window and search area sizes, and therefore the STARAN software is not suitable with the DIMP in the general case.

---

<sup>8</sup>D.J. Panton, M.E. Murphy, and D.S. Hanson, *Digital Cartographic Study and Benchmark; Final Report*, ETL-0168, U.S. Army Engineer Topographic Laboratories, Fort Belvoir, VA, 1978, AD-A064 800.

## CONCLUSIONS

1. The match point adjustment scheme used in the DIMP provides a highly successful means for correcting bad correlation results. Since the adjustment is applied at the completion of each row of correlations, the benefits are available immediately to minimize the occurrence of errors in subsequent rows of correlations.
2. The iterative logic technique is highly successful in preventing the correlation process from getting "lost," since it allows for a limited search operation when needed.
3. Window shaping generally improves correlation results, especially in steep terrain.
4. Match point prediction based on the track of previous correlation results is a necessity in steep terrain, especially where the slope of the terrain changes directions.
5. The linear correlation coefficient gives a more reliable and consistent measure of correlation than does the covariance or absolute difference coefficients.
6. The interactive techniques used in the DIMP provide a means for detecting and correcting correlation errors. The chance for further error propagation and the subsequent need for data editing are thereby minimized.
7. The visual 3-D display of correlation results as profiles on the ground gives the operator a reliable means for judging the accuracy of the correlation process.
8. The DIMP is a satisfactory laboratory tool for conducting tests involving digital correlation techniques. A special-purpose, hardwired computer is needed to perform the computation for the correlation coefficients before the DIMP can be made fast enough for production purposes.